Row Reduction of a Matrix and $A = c\,a\,r$

Steven L. Lee
slee@ornl.gov
www.ornl.gov/~slee

Gilbert Strang
gs@math.mit.edu
www-math.mit.edu/~gs

August 21, 1998

Abstract

Every matrix has a unique reduced row echelon form $R = \texttt{rref}(A)$. Algorithmically, the m by n matrix A is reduced to R by a sequence of elementary row operations. The product of those operations is an elimination matrix E such that $EA = R$. Part of E is uniquely determined and part depends on the sequence of steps. Similarly part of E^{-1} is determined and part is not. We show how E and R yield natural bases for the four fundamental subspaces (the column spaces and nullspaces of A and A^{T}). They also produce a factorization $A = C^* R^*$ into pivot columns times reduced rows.

This leads to the *echelon factorization* $A = c\,a\,r$ pointed out to us by Hans Schneider: r contains the nonzero rows of $\texttt{rref}(A)$, c contains the nonzero columns of $(\texttt{rref}(A'))'$, and a is the nonsingular submatrix formed by the pivot rows and columns of A. This factorization tells everything about the reduction to echelon bases for the row space and column space.

All these steps are conveniently described by single commands in our Linear Algebra Teaching Codes. The commands are repeated on the web page `http://web.mit.edu/18.06/www`. Students find them valuable in executing (and understanding!) the basic steps in the linear algebra course.

It is hard to say anything new about row reducing a matrix. Very likely this paper will not succeed. But in writing again about linear algebra (and creating short Teaching Codes for students in the course), we found some ideas that were new to us. We are always working with an m by n matrix A of rank r. We felt sure that the reduced row echelon form $R = \texttt{rref}(A)$ could make clear all four of the fundamental subspaces, and finally it did.

Our first example will be very small, just 2 by 2 of rank 1. If students can do the reduction by hand, so can we!

$$A = \begin{bmatrix} 2 & 6 \\ 1 & 3 \end{bmatrix} \quad \text{reduces to} \quad R = \begin{bmatrix} 1 & 3 \\ 0 & 0 \end{bmatrix}.$$

The reduced form has an r by r identity submatrix in the first r rows. All remaining nonzeros are to the right of those 1's, and R ends with $m - r$ zero rows. Our example has a 1 by 1 identity submatrix and one zero row. The first question is, how did we get from A to R?

Algebraically, R is uniquely determined by A. The "elementary row operations" always reach the same $EA = R$, but not always by the same matrix E. That is an important point for this paper. We want to understand E in case it is not unique.

In this example we might have multiplied row 1 by $1/2$, and then subtracted it from row 2. The result is R, and the product of those two elementary operations is E:

$$\begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1/2 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1/2 & 0 \\ -1/2 & 1 \end{bmatrix} = E.$$

Or we could exchange rows 1 and 2 of A, and then subtract 2 times row 1 from row 2:

$$\begin{bmatrix} 1 & 0 \\ -2 & 1 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & -2 \end{bmatrix} = E_{\text{new}}.$$

When we apply E or E_{new} (or some third possibility not yet identified) to A, the result is the same R:

$$EA = R \quad \text{and} \quad E_{\text{new}}A = R.$$

Those matrices E and R must contain all the crucial information about A, in an extremely convenient form. The matrix factorization $EA = R$ answers virtually every question that a linear algebra course traditionally asks! (Except eigenvalues.)

One case is clear from the start. If A is square and invertible, its reduced form must be $R = I$. Then the elimination matrix E is necessarily A^{-1}. The nullspaces of A and A^{T} contain only the zero vector—not very interesting. But when the rank r is less than m and n, we have four subspaces of fundamental importance:

The four subspaces are the column spaces and nullspaces of A and A^{T}.

Two subspaces are in $\mathbf{R}^n$ and two are in $\mathbf{R}^m$. The Fundamental Theorem gives the dimension of these subspaces. The column spaces share the same dimension r. The nullspaces have dimensions $n-r$ and $m-r$. In our example the subspaces are all one-dimensional, just lines in $\mathbf{R}^2$. The all-important picture in the linear algebra course shows the four fundamental subspaces and their perpendicularity:

PICTURE

Our first question is, *how do E and R reveal a basis for each subspace?* The r pivot columns are crucial. Those columns are <u>not</u> combinations of earlier columns; they are the first r independent columns of A. In the echelon form R, these are the columns that hold the r by r identity submatrix. MATLAB creates a list `pivcol` of the pivot column numbers. Then three of the four bases come directly from R:

1. Basis for the column space $C(A)$:
   the r columns of A listed in `pivcol`.

2. Basis for the row space $C(A^{\mathrm{T}})$:
   the first r rows of R.

3. Basis for the nullspace $N(A)$:
   the $n - r$ "special solutions" to $Rx = 0$ (and $Ax = 0$).

4. Where is a natural basis for the other nullspace $N(A^{\mathrm{T}})$?

We return to our example, as a linear algebra course should. The first column of R contains the pivot, so `pivcol` = [1]. Then the first column of A (not R!) is a basis for its column space. The columns of A and R are different, but *the column numbers in* `pivcol` *hold the key.*

The rows of A and R span the same row space. A basis is in the first r rows of R. And the solutions to $Ax = 0$ and $Rx = 0$ produce the same nullspace (this was the purpose of elimination). We will soon propose a basis of special solutions. First we answer the question that is still open, to identify the "left" nullspace $N(A^{\mathrm{T}})$.

This subspace is not revealed by R, but a basis is immediately visible in E:

The last $m - r$ rows of E are a basis for the left nullspace $N(A^{\mathrm{T}})$.

The reason is that $EA = R$ ends in $m - r$ zero rows. Those last rows of E are multiplying A (from the left) to produce the zero rows of R. Our example has one zero row in $EA = R$ and again in $E_{\text{new}}A = R$:

$$\begin{bmatrix} 1/2 & 0 \\ -1/2 & 1 \end{bmatrix}\begin{bmatrix} 2 & 6 \\ 1 & 3 \end{bmatrix} = \begin{bmatrix} 1 & 3 \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 0 & 1 \\ 1 & -2 \end{bmatrix}\begin{bmatrix} 2 & 6 \\ 1 & 3 \end{bmatrix} = \begin{bmatrix} 1 & 3 \\ 0 & 0 \end{bmatrix}.$$

Finally we notice: The last rows of E and E_{new} span the same space. This is the missing fourth subspace, the nullspace of A^{T}.

The Teaching Codes

Our starting point was to create Teaching Codes that would be helpful in the basic linear algebra course. Those codes are freely available from our web page. The commands were written to fit exactly with the Second Edition of the textbook *Introduction to Linear Algebra* [3]. The beauty of the codes is that they need only a few lines of MATLAB or Maple or Mathematica, and we hope it will be useful to include them here.

The command `rref` is available in all versions of MATLAB. The list `pivcol` is also available as a second output:

$$\begin{bmatrix} R, & \texttt{pivcol} \end{bmatrix} = \texttt{rref}(A).$$

Then our basis of pivot columns is immediate in the columns of

$$C^* = A(\ :\ ,\texttt{pivcol}).$$

The colon symbol before the comma produces all rows of the columns listed in `pivcol`. Similarly a colon after the comma produces all columns of the first r rows of R:

$$R^* = R(1 : \texttt{rank}(A),\ :\ ).$$

The special solutions to $Rx = 0$, whose discussion was postponed until now, will be the columns of the nullspace matrix N. The nullspace has dimension $n - r$, so N has $n - r$ columns. The nullspace matrix takes its simplest form when the r pivot columns come first in R:

$$RN = \begin{bmatrix} I & F \\ 0 & 0 \end{bmatrix} \begin{bmatrix} -F \\ I \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

Block multiplication is extremely useful!

In general `pivcol` must locate that r by r identity matrix in R. The Teaching Code is called `nulbasis`:

```
function N = nulbasis(A)
[R, pivcol] = rref(A);
[m, n] = size(A);
r = length(pivcol);
freecol = 1:n;
freecol(pivcol) = [ ];
N = zeros(n, n-r);
N(freecol, : ) = eye(n-r);
N(pivcol,  : ) = -R(1:r, freecol);
```

Always $RN = 0$! After the neat command involving the empty matrix `[ ]`, `freecol` lists the nonpivot columns. The free variable part of N is an identity matrix of order $n - r$. Then $Rx = 0$ determines the r pivot variables:

$$Rx = I \begin{bmatrix} \text{pivot} \\ \text{variables} \end{bmatrix} + F \begin{bmatrix} \text{free} \\ \text{variables} \end{bmatrix} = 0.$$

So the pivot variables are given by $-F$ in the last line of the code.

Example: Suppose R is 2 by 4 with `pivcol` = [1 3]. Then N has $4 - 2$ columns:

$$R = \begin{bmatrix} 1 & 2 & 0 & 3 \\ 0 & 0 & 1 & 4 \end{bmatrix} \text{yields } N = \begin{bmatrix} -2 & -3 \\ 1 & 0 \\ 0 & -4 \\ 0 & 1 \end{bmatrix}.$$

This nulbasis construction is so simple, but neat. The nullspace matrix N is in "reverse echelon form." Its extra nonzeros are *above the ones* in the identity submatrix. The list `freecol` for R, which is [2 4], gives the row numbers of that identity matrix in N.

The row space has a unique echelon basis, and its orthogonal complement the nullspace has a unique reverse echelon basis. The two are connected by `nulbasis`.

Finally we come to E, the product of all the elementary elimination steps. Its last $m - r$ rows form a basis for the left nullspace. The easy way to construct E follows the Gauss-Jordan idea of augmenting A to $\begin{bmatrix} A & I \end{bmatrix}$. *Then row reduction produces* $\begin{bmatrix} R & E \end{bmatrix}$. The command is $E =$ `elim`(A) and the Teaching Code separates R from E:

```
function [E, R] = elim(A)
[m, n] = size(A);
I = eye(m);
RE = rref([A I]);
R = RE(:, 1:n);
E = RE(:, (n+1):(m+n));
```

The operations that change A into R will at the same time change I into E. The long matrix $\begin{bmatrix} A & I \end{bmatrix}$ is just multiplied by the elimination matrix:

$$E \text{ multiplies } \begin{bmatrix} A & I \end{bmatrix} \text{ to give } \begin{bmatrix} R & E \end{bmatrix}.$$

In the square invertible case $EA = R$ becomes $EA = I$. This is how the Gauss-Jordan method finds $E = A^{-1}$ (quite efficiently).

Only one more question to go. *Is it* E *or* E_{new} *or some third matrix that is produced by* `elim`(A)*?* This is a definite command, and it is going to produce a definite E. How can we tell which one, when a whole family of E's satisfies $EA = R$? We try the example first:

$$\begin{bmatrix} A & I \end{bmatrix} = \begin{bmatrix} 2 & 6 & 1 & 0 \\ 1 & 3 & 0 & 1 \end{bmatrix} \text{ reduces to } \begin{bmatrix} R & E \end{bmatrix} = \begin{bmatrix} 1 & 3 & 0 & 1 \\ 0 & 0 & 1 & -2 \end{bmatrix}.$$

The Teaching Code `elim` has picked out E_{new} in those last two columns. Why this choice? Because $\begin{bmatrix} R & E \end{bmatrix}$ always has a full-size identity matrix in the pivot columns of $\begin{bmatrix} A & I \end{bmatrix}$.

Those pivot columns are uniquely determined! The command `elim` must have chosen this particular elimination matrix because E_{new}^{-1} has all the pivot columns of $\begin{bmatrix} A & I \end{bmatrix}$:

$$E_{\text{new}}^{-1} = \begin{bmatrix} 2 & 1 \\ 1 & 0 \end{bmatrix} \quad \text{contains columns 1 and 3 of} \quad \begin{bmatrix} 2 & 6 & 1 & 0 \\ 1 & 3 & 0 & 1 \end{bmatrix}.$$

The Factorizations $A = E^{-1}R$ and $A = C^*R^*$

The great factorizations of linear algebra are traditionally written in the form $A = (\ )(\ )$ or $A = (\ )(\ )(\ )$. The right sides contain triangular matrices or diagonal matrices or orthogonal matrices. The most famous is $A = LU$ with elimination producing the triangular factors. Close behind comes $A = QR =$ (orthogonal)(triangular) from the Gram-Schmidt process (or Householder reduction). Somehow $A = E^{-1}R$ must be one more matrix factorization with its own special features.

We already know what is special about the echelon form R. But E^{-1} is still a mystery. All we have discovered is that the last $m - r$ rows of E are a basis for the nullspace of A^{T}. *If we know something about rows of E, what does that reveal about E^{-1}?* If we know only the last row of E, then *all but the last column of E^{-1} would be perpendicular to it.* This fact is expressed by the last row of $EE^{-1} = I$. And, more generally:

The first r columns of E^{-1} are perpendicular to the last $m - r$ rows of E.

In our situation, those last $m-r$ rows produce the nullspace of A^{T}. Therefore the first r columns of E^{-1} must produce the perpendicular subspace. *The Fundamental Theorem says that this is the column space of A.* Look at the first column in our example:

$$E^{-1} = \begin{bmatrix} 2 & 0 \\ 1 & 1 \end{bmatrix} \quad \text{and} \quad E_{\text{new}}^{-1} = \begin{bmatrix} 2 & 1 \\ 1 & 0 \end{bmatrix}.$$

That first column (shared by E^{-1} and E_{new}^{-1}) is a basis for the column space of A.

More than that, the first column is exactly the pivot column of A. It is the special basis we identified using `pivcol`. There must be a reason for this particular first column, and we see it from the factorizations $E^{-1}R$ and $E_{\text{new}}^{-1}R$:

$$A = \begin{bmatrix} 2 & 0 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 3 \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} 2 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 3 \\ 0 & 0 \end{bmatrix}.$$

The first r columns of E^{-1} are multiplying the *identity matrix* in R to give the pivot columns of A. So those first columns are precisely the pivot columns, and we have a useful factorization:

Theorem 1. Every m by n matrix of rank r can be factored into $A = E^{-1}R$, where $R = \texttt{rref}(A)$ and the first r columns of E^{-1} are the pivot columns of A. The last $m - r$ columns of E^{-1} are not determined because they multiply the zero rows of R.

Let us illustrate this factorization when the pivot columns happen to come first:

$$A = E^{-1}R = \begin{bmatrix} C^* & Z \end{bmatrix} \begin{bmatrix} I & F \\ 0 & 0 \end{bmatrix} = C^* \begin{bmatrix} I & F \end{bmatrix} + Z \begin{bmatrix} 0 & 0 \end{bmatrix}.$$

C^* contains all pivot columns of A and $\begin{bmatrix} I & F \end{bmatrix}$ contains all nonzero rows of R. But Z is still undetermined. Seeing those undetermined columns of E^{-1} multiplying those zero rows of R, why not just remove them? This leaves a canonical factorization of A into an m by r matrix C^* (r columns) and an r by n matrix R^* (r rows):

Corollary. $$A = C^*R^* = \begin{bmatrix} \text{pivot} \\ \text{columns} \\ \text{of } A \end{bmatrix} \begin{bmatrix} \text{nonzero rows} \\ \text{of } R \end{bmatrix}. \qquad (1)$$

These factors contain our favorite bases for the column space and the row space of A. There is no assumption that the pivot columns come first. This C^*R^* factorization is "known but not well known." Let us describe the factorization in simple words: C^* contains the pivot columns of A, and R^* tells how each column of A is a combination of those pivot columns.

The Factorization $A = c\,a\,r$

Hans Schneider pointed out a more symmetric factorization that really is canonical [2]. It has three factors instead of two. So far we have a reduced echelon basis for the row space of A (in the rows of R^*). It is natural to ask for a reduced echelon basis for the column space too. This gives the (possibly new?) factorization in the title of our paper.

In between the two echelon matrices comes the leading r by r nonsingular submatrix of A. We call it a.

Theorem 2. Every m by n matrix of rank r factors into

$$A = c\,a\,r = (m \text{ by } r)(r \text{ by } r)(r \text{ by } n). \tag{2}$$

The matrix $r = R^*$ contains the nonzero rows of `rref`(A). The matrix c contains the nonzero columns of (`rref`$(A'))'$. The matrix a is the nonsingular submatrix formed by the pivot columns of A and its pivot rows (the pivot columns of A^{T}).

The algorithm below yields one proof, and the reasoning is equally quick in English. Start from the C^*R^* factorization, with the pivot columns of A in C^*. Then factor C^* into $c\,a$. This says that each row of C^* is a combination of its pivot rows (which are the rows of a). The echelon matrix c executes those linear combinations of the pivot rows.

Actually we are using equation (1) to produce C^*R^*, and then using it again (transposed) to factor C^* into $c\,a$. With $R^* = r$ we have $A = c\,a\,r$.

Thus r and c contain the unique "reduced echelon bases" for the row space and column space. Both include an r by r identity submatrix I. All other nonzeros are below I (in c) and to the right (in r). In case I comes first in both matrices, the $c\,a\,r$ factorization is simply

$$A = \begin{bmatrix} I \\ G \end{bmatrix} \begin{bmatrix} a \end{bmatrix} \begin{bmatrix} I & F \end{bmatrix},$$

and the nonsingular submatrix a is in the upper left corner of A. In general it is in the pivot columns and pivot rows. A small theorem is concealed here: The list `pivcol` stays the same after premultiplication by an invertible matrix.

A short Echelon Factorization code produces c and a and r:

```
function [c, a, r] = car(A)
[R, pivcol] = rref(A);
[S, pivrow] = rref(A');
r = R(1:rank(A), : );
c = S(1:rank(A), : )';
a = A(pivrow, pivcol);
```

Example of $c\,a\,r$ with `pivcol` $= [1\ 3]$ and `pivrow` $= [1\ 2]$:

$$A = \begin{bmatrix} \mathbf{1} & 1 & \mathbf{1} & 1 \\ \mathbf{1} & 1 & \mathbf{2} & 3 \\ 2 & 2 & 3 & 4 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{1} & \mathbf{1} \\ \mathbf{1} & \mathbf{2} \end{bmatrix} \begin{bmatrix} 1 & 1 & 0 & -1 \\ 0 & 0 & 1 & 2 \end{bmatrix}.$$

The submatrix A has full rank and is as far "northwest" as possible in A. It is the leading r by r invertible submatrix, because `pivrow` and `pivcol` pick out the first r independent rows and columns.

One final point. Every matrix B with the same column space and row space as A must have the form $B = c\,b\,r$. The r by r invertible matrix in the middle parameterizes all matrices that share the same four fundamental subspaces. And when A is square and invertible, $c = I$ and $r = I$ and a is A!

Historical note. It is interesting to ask where the factors should go. Shall we put an invertible P and Q on the left side, multiplying A? Or shall the factors go on the right side as in $A = c\,a\,r$? The classical theorem, neatly discussed by Cohn [1, pg. 61], chooses P and Q to reveal the rank as the only invariant:

$$PAQ = \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix}.$$

We like the modern form because c, a, and r inherit direct meanings from A. A more famous example compares $A = LU$ with $L^{-1}A = U$. The entries of L are the multipliers in elimination—they have individual meanings. By contrast, L^{-1} is an important matrix (it triangularizes A) but its entries are obscure.

Row Reduction for Block Matrices

This brief section mentions some exercises that help students to understand R. After practicing on ordinary matrices A, we ask about a few *block matri-*

ces. Probably the questions will be evident from their answers:

1. `rref` applied to $\begin{bmatrix} A & A \end{bmatrix}$ produces $\begin{bmatrix} R & R \end{bmatrix}$.

2. `rref` applied to $\begin{bmatrix} A \\ A \end{bmatrix}$ produces $\begin{bmatrix} R \\ 0 \end{bmatrix}$. (What is E?)

3. `rref` applied to $\begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix}$ produces $\begin{bmatrix} R & 0 \\ 0 & R \end{bmatrix}$.

Not quite! The zero rows in the upper R must move to the bottom to produce $\begin{bmatrix} R^* & 0 \\ 0 & R^* \\ 0 & 0 \end{bmatrix}$.

4. `rref` applied to $\begin{bmatrix} A & A \\ A & 0 \end{bmatrix}$ produces which reduced form?

We can certainly reach $\begin{bmatrix} R & R \\ R & 0 \end{bmatrix}$ and then $\begin{bmatrix} R & R \\ 0 & -R \end{bmatrix}$ and then $\begin{bmatrix} R & 0 \\ 0 & R \end{bmatrix}$. Again all zero rows must move to the bottom.

What is the principle that governs these examples? They are all tensor products of small matrices B with A. The small matrices in Examples 1 to 4 are just

$$B = \begin{bmatrix} 1 & 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 \\ 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}.$$

The reduced row echelon forms of those matrices B are

$$\texttt{rref}(B) = \begin{bmatrix} 1 & 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

It looks as if the rule is this: *The reduced form of* $B \otimes A$ *is* $\texttt{rref}(B) \otimes \texttt{rref}(A)$*, except that all zero rows move to the bottom.*

It is interesting that adding or multiplying matrices makes a mess of their row reduced forms, while the block operation of tensor products is quite reasonable.

Sums and Intersections

The same commands can produce bases for the sum and intersection of two subspaces of $\mathbf{R}^m$. Suppose the columns of A are a basis for S and the columns of B are a basis for T. Then the columns of $\begin{bmatrix} A & B \end{bmatrix}$ span the sum $S+T$. The command

$$\texttt{sumbasis} = \texttt{colbasis}(\begin{bmatrix} A & B \end{bmatrix})$$

quickly picks out a basis for $S+T$. Its dimension is the rank of $\begin{bmatrix} A & B \end{bmatrix}$.

The intersection $S \cap T$ is not so obvious. One expects it to come from the nullspace of $\begin{bmatrix} A & B \end{bmatrix}$. If so, the connection formula

$$\texttt{dim}(S+T) + \texttt{dim}(S \cap T) = \texttt{dim}(S) + \texttt{dim}(T) \tag{3}$$

is immediate from the Fundamental Theorem. The right side is the number of columns in $\begin{bmatrix} A & B \end{bmatrix}$. On the left side is the dimension of the column space and (we hope) the nullspace. But how do we identify $S \cap T$ in $\mathbf{R}^m$ with the nullspace of $\begin{bmatrix} A & B \end{bmatrix}$?

In our hardest problem set, we asked the linear algebra class to create simple examples of $S+T$ and $S \cap T$. From their examples they could guess formula (3). Most of them came up with the command $\texttt{colbasis}(\begin{bmatrix} A & B \end{bmatrix})$ to deal with the sum $S+T$ (as above). But the difficult question—to find a basis for $S \cap T$ from the matrices A and B—separated out the best from the also-rans. The instructors found themselves among the also-rans, as we now humbly explain.

Our idea for $S \cap T$ came from a recent paper by Yang [4]. The second author remembered a similar suggestion in his earlier book: each vector in the nullspace of $\begin{bmatrix} A & B \end{bmatrix}$ combines the columns to give the zero vector. So a combination of columns of A (a vector in S) equals a combination of columns of B (a vector in T). This locates a vector in $S \cap T$. The problem is to verify that this map from the nullspace of $\begin{bmatrix} A & B \end{bmatrix}$ to $S \cap T$ is an isomorphism, so the dimensions match:

$$\texttt{dimension of nullspace} = \texttt{dimension of } S \cap T.$$

Our class had never heard of an isomorphism, so this was a rough road.

Two students came forward after class with another idea. They knew that the intersection of two sets is the complement of the union of the complements:

$$S \cap T = (S^c \cup T^c)^c. \tag{4}$$

So if "union" is replaced by "sum" and the complement of a set changes to the orthogonal complement of a subspace, they had what they wanted. And they knew the Fundamental Theorem: `nulbasis`(A') produces a complement to `colbasis`(A). So the set identity (4) converted into a basis formula for the intersection of subspaces:

$$\texttt{intbasis} = \texttt{nulbasis}\left(\begin{bmatrix} \texttt{nulbasis}(A') & \texttt{nulbasis}(B') \end{bmatrix}'\right) \tag{5}$$

Their names were Yan and Dianne. They got A's on the spot.

References

[1] P.M. Cohn, *Elements of Linear Algebra*, Chapman and Hall, London, 1994.

[2] Hans Schneider, *Personal communication.*

[3] Gilbert Strang, *Introduction to Linear Algebra, Second Edition*, Wellesley-Cambridge Press, Wellesley MA, 1998.

[4] K.-W. Yang, *A Basis for the Intersection of Subspaces*, Mathematics Magazine, Vol. 70, October 1997, pg. 297.